

Linear Credit Risk Models ^{*}

Damien Akerer [†] Damir Filipović [‡]

May 19, 2016

Abstract

We introduce a novel class of credit risk models in which the drift of the survival process of a firm is a linear function of the factors. These models outperform the standard affine default intensity models in terms of analytical tractability. The prices of defaultable bonds and credit default swaps (CDS) are linear in the factors. The price of a CDS option can be uniformly approximated by polynomials in the factors. An empirical study illustrates the versatility of these models by fitting CDS spread time series.

Keywords: Credit Default Swap, Credit Default Swap Option, Credit Risk, Credit Valuation Adjustment, Survival Process.

JEL Classification: G12, G13

^{*}The authors would like to thank for useful comments Agostino Capponi and Martin Larsson, as well as participants from the 2015 AMaMeF and Swissquote conference in Lausanne. The research leading to these results has received funding from the European Research Council under the European Union's Seventh Framework Programme (FP/2007-2013) / ERC Grant Agreement n. 307465-POLYTE.

[†]EPFL and Swiss Finance Institute. Email: damien.ackerer@epfl.ch

[‡]EPFL and Swiss Finance Institute. Email: damir.filipovic@epfl.ch

1 Introduction

Credit risk is virtually inherent to all financial securities. Breach of contracts caused by the nonpayment of cash flows, as well as variations in asset values caused by changing default risk, are omnipresent in financial markets. The underestimation of credit risk before the financial crisis incited regulators around the globe to force financial institutions to better manage and report credit risk. The complexity of credit risky portfolios and the securities therein renders the valuation of credit risk a challenging task that calls for suitable models.

We introduce a novel class of flexible and tractable reduced form models for the term structure of credit risk, the linear credit risk models. We directly specify the survival process of a firm, that is, its conditional survival probability given the economic background information. Specifically, we assume a multivariate factor process with a linear drift and let the drift of the survival process be linear in the factors. Prices of defaultable bonds and credit default swaps (CDS) are given in closed-form as linear functions of the factors. The implied default intensity is an explicit linear-rational function of the factors. In contrast, the price of a CDS in an affine default intensity model is a sum of exponential-affine functions in terms of solutions of nonlinear ordinary differential equations that are not explicit. Albeit we focus on single-firm credit risk in this paper, the linear credit risk models carry over to multi-firm credit portfolios.

Within the linear framework we define the linear hypercube (LHC) model where the factor process is diffusive with quadratic diffusion function. The factor process takes values in a hypercube whose edges' length is given by the survival process. The quadratic diffusion function is concave and bi-monotonic. This feature allows factors to virtually jump between low and high values. This facilitates the persistence and likelihood of term structure shifts. The factors' volatility parameters do not enter the bond and CDS pricing formulas, yet they impact the volatility of CDS spreads and thus affect CDS options. This facilitates the fitting of time series. We discuss in detail the one-factor LHC model and compare it with the one-factor affine default intensity model. We provide an identifiable canonical representation and the market price of risk specifications that preserve the linear drift of the factors.

We present a price approximation methodology for European style options on credit risky underlyings that exploits the compactness of the state space

and the closed-form of the conditional moments of the factor process. First, by the Stone–Weierstrass theorem, any continuous payoff function on the compact state space can be approximated by a polynomial to any given level of accuracy. Second, the conditional expectation of any polynomial in the factors is a polynomial in the prevailing factor values. In consequence, the price of a CDS option can be uniformly approximated by polynomials in the factors. This method also applies to the computation of credit valuation adjustments.

We perform an empirical analysis of the LHC model. Assuming a parsimonious cascading drift structure, we fit two- and three-factor LHC models to the ten-year long time series of weekly CDS spreads on a major US bank. The three-factor model is able to capture the complex term structure dynamics remarkably well and performs significantly better than the two-factor model. To illustrate the efficiency of the option pricing method we approximate the prices of CDS options with different moneyness. Polynomials of relatively low orders are sufficient to obtain accurate proxies for in-the-money options. Out-of-the money options typically require a higher order. We also fit some affine default intensity models and show that they have a similar ability in fitting the CDS spreads.

We now review some of the related literature. Our approach follows a standard doubly stochastic construction of default times as described in (Elliott, Jeanblanc, and Yor 2000). The early contributions by (Lando 1998) and (Duffie and Singleton 1999) already make use of affine factor processes. In contrast, the factor process in the LHC model is a strictly non-affine polynomial diffusion, whose general properties are studied in (Filipović and Larsson 2016). The stochastic volatility models developed in (Hull and White 1987) and (Akerer, Filipović, and Pulido 2016) are two other examples of non-affine models. Factors in the LHC models have a compact support and can exhibit jump-like dynamics similar to the multivariate Jacobi process introduced by (Gourieroux and Jasiak 2006). Our approach has some similarities with the linearity generating process by (Gabaix 2009) and the linear-rational models by (Filipović, Larsson, and Trolle 2016). These models exploit the tractability of the factor processes with linear drift and target non default risky asset pricing. To the authors knowledge, our paper is the first to model directly the survival process of a firm with linear drift characteristics.

Options on CDS contracts are complex derivatives and intricate to price. The pricing and hedging of CDS options in a generic hazard process framework is discussed in (Bielecki, Jeanblanc, and Rutkowski 2006) and (Bielecki,

Jeanblanc, Rutkowski, et al. 2008), and specialised to the square-root diffusion factor process in (Bielecki, Jeanblanc, and Rutkowski 2011). More recently (Brigo and El-Bachir 2010) developed a semi-analytical expression for CDS option prices in the context of a shifted square-root jump-diffusion default intensity model that was introduced in (Brigo and Alfonsi 2005). Another strand of the literature has focused on developing market models in the spirit of LIBOR market models. We refer the interested reader to (Schönbucher 2000), (Hull and White 2003), (Schönbucher 2004), (Jamshidian 2004), and (Brigo and Morini 2005). Although offering more tractability, this approach makes it difficult, if not impossible, to consistently value multiple instruments exposed to the same credit risk. The idea of approximating option prices by power series can be traced back to (Jarrow and Rudd 1982). However, most of the previous literature has focussed on approximating the transition density function of the underlying process, see for example (Corrado and Su 1996) and (Filipović, Mayerhofer, and Schneider 2013). In contrast, we approximate directly the payoff function by a power series.

The remainder of the paper is structured as follows. Section 2 presents the linear credit risk framework. Section 3 describes the LHC model. The empirical analysis is in Section 4. Section 5 concludes. The appendix contains all proofs and some additional results on market price of risk specifications that preserve the linear drift of the factors.

2 The Linear Framework

We introduce the linear credit risk model framework and derive closed-form expressions for defaultable bond prices and CDS. We also discuss CDS option pricing and credit valuation adjustments.

2.1 Survival Process Specification

We fix a stochastic basis $(\Omega, \mathcal{F}, \mathcal{F}_t, \mathbb{Q})$, where \mathcal{F}_t represents the economic background information and \mathbb{Q} is the risk-neutral pricing measure. Let S_t be the survival process of a firm. This is a right-continuous \mathcal{F}_t -adapted and non-increasing process such that $0 < S_t \leq 1$. Let U be a standard uniform random variable which is independent from \mathcal{F}_∞ . We define the random time

$$\tau := \inf\{t \geq 0 \mid S_t \leq U\},$$

which is infinity if the set is empty. Let \mathcal{H}_t be the filtration generated by the indicator process $H_t = \mathbb{1}_{\{\tau > t\}}$, which is one as long as the firm has not defaulted by time t , and zero afterwards. The default time τ is a stopping time in the enlarged filtration $\mathcal{G}_t = \mathcal{F}_t \vee \mathcal{H}_t$. It is \mathcal{F}_t -doubly stochastic in the sense that

$$\mathbb{Q}[\tau > t \mid \mathcal{F}_\infty] = \mathbb{Q}[S_t > U \mid \mathcal{F}_\infty] = S_t.$$

In a linear credit risk model the dynamics of S_t is of the form

$$dS_t = -\gamma^\top X_t dt - dM_t^S$$

for some $\gamma \in \mathbb{R}^m$, some \mathcal{F}_t -martingale M_t^S , and some m -dimensional factor process X_t with linear drift of the form

$$dX_t = (\beta S_t + B X_t) dt + dM_t^X$$

for some $\beta \in \mathbb{R}^m$, $B \in \mathbb{R}^{m \times m}$, and m -dimensional \mathcal{F}_t -martingale M_t^X . The linear drift of (S_t, X_t) implies that the \mathcal{F}_t -conditional expectation of (S_T, X_T) is linear of the form

$$\mathbb{E} \left[\begin{pmatrix} S_T \\ X_T \end{pmatrix} \mid \mathcal{F}_t \right] = e^{A(T-t)} \begin{pmatrix} S_t \\ X_t \end{pmatrix}, \quad t \leq T, \quad (1)$$

where the $(1+m) \times (1+m)$ -matrix A is defined as

$$A = \begin{pmatrix} 0 & -\gamma^\top \\ \beta & B \end{pmatrix}.$$

Remark 2.1. *If S_t is absolutely continuous, such that $dM_t^S = 0$ for all $t \geq 0$, the corresponding default intensity λ_t that derives from the relation $S_t = e^{-\int_0^t \lambda_s ds}$ is linear-rational of the form*

$$\lambda_t = \frac{\gamma^\top X_t}{S_t}.$$

2.2 Defaultable Bonds

We assume a constant risk-free interest rate equal to r such that the time- t price of the risk-free zero-coupon bond price with maturity T and face value one is given by $e^{-r(T-t)}$. The following result gives a closed-form expression for the price of a defaultable bond with constant recovery rate at maturity.

Proposition 2.2. *The time- t price of a defaultable zero-coupon bond with maturity T and recovery δ at **Maturity** is*

$$B^M(t, T) = (1 - \delta)B^Z(t, T) + \mathbb{1}_{\{\tau > t\}}\delta e^{-r(T-t)}$$

where

$$B^Z(t, T) = e^{-r(T-t)}\mathbb{E} [\mathbb{1}_{\{\tau > T\}} | \mathcal{G}_t]$$

denotes the time- t price of a defaultable zero-coupon bond with maturity T and **Zero** recovery. It is of the form

$$B^Z(t, T) = \mathbb{1}_{\{\tau > t\}} \frac{e^{-r(T-t)}}{S_t} \psi_Z(t, T)^\top \begin{pmatrix} S_t \\ X_t \end{pmatrix}$$

where the vector $\psi_Z(t, T) \in \mathbb{R}^{1+m}$ is given by

$$\psi_Z(t, T)^\top = \mathbf{e}_1^\top e^{A(T-t)}$$

and $\mathbf{e}_1 \in \mathbb{R}^{1+m}$ denotes the vector whose first coordinate is equal to one and the others are equal to zero.

The next result shows that also the price of a defaultable bond paying a constant recovery rate at default can be retrieved in closed-form.

Proposition 2.3. *The time- t price of a defaultable zero-coupon bond with maturity T and recovery δ at **Default** is*

$$B^D(t, T) = B^Z(t, T) + \delta C^D(t, T),$$

where

$$C^D(t, T) = \mathbb{E} [e^{-r(\tau-t)} \mathbb{1}_{\{t < \tau \leq T\}} | \mathcal{G}_t]$$

denotes the time- t price of a contingent cash flow paying one at default if it occurs between dates t and T . It is of the form

$$C^D(t, T) = \mathbb{1}_{\{\tau > t\}} \frac{1}{S_t} \psi_D(t, T)^\top \begin{pmatrix} S_t \\ X_t \end{pmatrix}$$

where the vector $\psi_D(t, T) \in \mathbb{R}^{1+m}$ is given by

$$\psi_D(t, T)^\top = (0 \quad \gamma^\top) A_*^{-1} (e^{A_*(T-t)} - \mathbf{I})$$

with $A_* = A - \mathbf{I}r$ and \mathbf{I} denotes the identity matrix of dimension $1 + m$.

Proposition 2.3 is a noticeable result as the pricing of such contingent cash flows usually involves numerical integration. For illustration assume that the survival process S_t is absolutely continuous admitting the default intensity λ_t as in Remark 2.1. Then $C^D(t, T)$ can be rewritten as

$$C^D(t, T) = \mathbb{1}_{\{\tau > t\}} \int_t^T e^{-r(u-t)} \mathbb{E} \left[\lambda_u e^{-\int_t^u \lambda_s ds} \mid \mathcal{F}_t \right] du.$$

With affine default intensity models the expectation to be integrated requires solving Riccati equations which typically do not have an explicit solution. Numerical methods such as finite difference are usually employed to compute the expectation with time- u cash flow for $u \in [t, T]$. The integral can then only be approximated by means of another numerical method such as quadrature, that necessitates solving the corresponding ODEs at many different points u . The pricing formulas for affine jump-diffusion models can be found in (Duffie, Pan, and Singleton 2000). By avoiding those operations the linear framework therefore is computationally more efficient and avoid numerical accuracy issues.

The price of a security whose only cash flow is proportional to the default time is given in the following lemma. It is of interest to compute the expected accrued interests at default for some contingent securities such as CDS.

Lemma 2.4. *The time- t price of a contingent bond paying τ at default if it occurs between date t and T is of the form*

$$C^{D\tau}(t, T) = \mathbb{E} \left[\tau e^{-r(\tau-t)} \mathbb{1}_{\{t < \tau \leq T\}} \mid \mathcal{G}_t \right] = \mathbb{1}_{\{\tau > t\}} \frac{1}{S_t} \psi_{D\tau}(t, T)^\top \begin{pmatrix} S_t \\ X_t \end{pmatrix}$$

where the vector $\psi_{D\tau}(t, T) \in \mathbb{R}^{1+m}$ is given by

$$\psi_{D\tau}(t, T)^\top = \begin{pmatrix} 0 & \gamma^\top \end{pmatrix} \left((T-t) A_*^{-1} e^{A_*(T-t)} + A_*^{-1} (I t - A_*^{-1}) (e^{A_*(T-t)} - I) \right).$$

Similar to what we discussed after Proposition 2.3, the explicit formula given in Lemma 2.4 avoids the numerical integration of the present value of the cash flow $u e^{-r(u-t)} \mathbb{E} \left[\lambda_u e^{-\int_t^u \lambda_s ds} \mid \mathcal{F}_t \right]$ over the period $[t, T]$.

2.3 CDS and CDS Options

A CDS is an insurance contract that pays at default the realized loss on a reference bond – the protection leg – in exchange of periodic payments that

will stop after default – the premium leg. We consider the following discrete tenor structure $t \leq T_0 < T_1 < \dots < T_N$ and a contract offering default protection from date T_0 to date T_N . When $t < T_0$ the contract is usually called a knock out forward CDS and generates cash flows only if the firm has not defaulted by time T_0 . The time- t value of the premium leg with spread k is given by $k V_{\text{prem}}(t, T_0, T_N)$ where

$$V_{\text{prem}}(t, T_0, T_N) = V_{\text{coup}}(t, T_0, T_N) + V_{\text{ai}}(t, T_0, T_N)$$

is the sum of the value of coupon payments before default

$$V_{\text{coup}}(t, T_0, T_N) = \mathbb{E} \left[\sum_{k=1}^N e^{-r(T_k - t)} (T_k - T_{k-1}) \mathbb{1}_{\{T_k < \tau\}} \mid \mathcal{G}_t \right]$$

and the value of the accrued coupon payment at the time of default

$$V_{\text{ai}}(t, T_0, T_N) = \mathbb{E} \left[\sum_{k=1}^N e^{-r(\tau - t)} (\tau - T_{k-1}) \mathbb{1}_{\{T_{k-1} < \tau \leq T_k\}} \mid \mathcal{G}_t \right].$$

The time- t value of the protection leg is

$$V_{\text{prot}}(t, T_0, T_N) = (1 - \delta) \mathbb{E} \left[e^{-r(\tau - t)} \mathbb{1}_{\{T_0 < \tau \leq T_N\}} \mid \mathcal{G}_t \right],$$

where δ denotes the constant recovery rate at default. The (forward) CDS spread $\text{CDS}(t, T_0, T_N)$ is the spread k that makes the premium leg and the protection leg equal in value at time t . That is,

$$\text{CDS}(t, T_0, T_N) = \frac{V_{\text{prot}}(t, T_0, T_N)}{V_{\text{prem}}(t, T_0, T_N)}.$$

Proposition 2.5. *The elements of the CDS spread $\text{CDS}(t, T)$ are of the form*

$$V_{\text{prot}}(t, T_0, T_N) = \mathbb{1}_{\{\tau > t\}} \frac{1}{S_t} \psi_{\text{prot}}(t, T_0, T_N)^\top \begin{pmatrix} S_t \\ X_t \end{pmatrix}$$

$$V_{\text{prem}}(t, T_0, T_N) = \mathbb{1}_{\{\tau > t\}} \frac{1}{S_t} \psi_{\text{prem}}(t, T_0, T_N)^\top \begin{pmatrix} S_t \\ X_t \end{pmatrix}$$

where the vectors $\psi_{\text{prot}}(t, T_0, T_N), \psi_{\text{prem}}(t, T_0, T_N) \in \mathbb{R}^{1+m}$ are given by

$$\begin{aligned}\psi_{\text{prot}}(t, T_0, T_N) &= (1 - \delta) (\psi_D(t, T_N) - \psi_D(t, T_0)), \\ \psi_{\text{prem}}(t, T_0, T_N) &= \sum_{k=1}^N (T_k - T_{k-1}) \psi_Z(t, T_k) + \psi_{D\tau}(t, T_N) - \psi_{D\tau}(t, T_0) \\ &\quad + T_{N-1} \psi_D(t, T_N) - \sum_{k=1}^{N-1} (T_k - T_{k-1}) \psi_D(t, T_k) - T_0 \psi_D(t, T_0).\end{aligned}$$

Proposition 2.5 implies that the CDS spread simplifies to a readily available linear-rational expression,

$$\text{CDS}(t, T_0, T_N) = \mathbb{1}_{\{\tau > t\}} \frac{\psi_{\text{prot}}(t, T_0, T_N)^\top \begin{pmatrix} S_t \\ X_t \end{pmatrix}}{\psi_{\text{prem}}(t, T_0, T_N)^\top \begin{pmatrix} S_t \\ X_t \end{pmatrix}}.$$

For comparison, in an affine default intensity model the two legs $V_{\text{prot}}(t, T_0, T_N)$ and $V_{\text{prem}}(t, T_0, T_N)$ are given as sums of exponential-affine terms that cannot be simplified further and where each term has to be computed numerically.

A CDS option with strike spread k is a European call option on the CDS contract exercisable only if the firm has not defaulted before the option maturity date T_0 . Its payoff is

$$\mathbb{1}_{\{\tau > T_0\}} (V_{\text{prot}}(T_0, T_0, T_N) - k V_{\text{prem}}(T_0, T_0, T_N))^+ = \mathbb{1}_{\{\tau > T_0\}} \frac{Z(T_0, T_N, k)^+}{S_{T_0}}$$

for the random variable

$$Z(T_0, T_N, k) = (\psi_{\text{prot}}(T_0, T_0, T_N) - k \psi_{\text{prem}}(T_0, T_0, T_N))^\top \begin{pmatrix} S_{T_0} \\ X_{T_0} \end{pmatrix}.$$

Denote $\text{CDSO}(t, T_0, T_N, k)$ the price of the CDS option at time t ,

$$\begin{aligned}\text{CDSO}(t, T_0, T_N, k) &= \mathbb{E} \left[e^{-r(T_0-t)} \mathbb{1}_{\{\tau > T_0\}} \frac{Z(T_0, T_N, k)^+}{S_{T_0}} \mid \mathcal{G}_t \right] \\ &= \mathbb{1}_{\{\tau > t\}} \frac{e^{-r(T_0-t)}}{S_t} \mathbb{E} [Z(T_0, T_N, k)^+ \mid \mathcal{F}_t].\end{aligned}\tag{2}$$

The price of a CDS option is therefore equal to the expected positive part of a linear function of (S_{T_0}, X_{T_0}) , adjusted for time value and realized

credit risk. When the characteristic function of the process (S_{T_0}, X_{T_0}) is available such expression can be computed efficiently using Fourier transform techniques. The next section presents a new methodology based on payoff approximation requiring only the knowledge of a finite number of conditional moments of (S_{T_0}, X_{T_0}) .

2.4 Credit Valuation Adjustment

The unilateral credit valuation adjustment (UCVA) of a position is the present value of losses resulting from its cancellation when a bilateral counterparty defaults.

Proposition 2.6. *The time- t price of the UCVA with maturity T and time- u net positive exposure $f(u, S_u, X_u)$, for some continuous function $f(u, s, x)$, is*

$$\begin{aligned} \text{UCVA}(t, T) &= \mathbb{E} \left[e^{-r(\tau-t)} \mathbb{1}_{\{t < \tau \leq T\}} f(\tau, S_\tau, X_\tau) \mid \mathcal{G}_t \right] \\ &= \mathbb{1}_{\{\tau > t\}} \frac{1}{S_t} \int_t^T e^{-r(u-t)} \mathbb{E} \left[f(u, S_u, X_u) \gamma^\top X_u \mid \mathcal{F}_t \right] du. \end{aligned}$$

where τ is the counterparty default time.

Computing the UCVA therefore boils down to a numerical integration of European style option prices. As is the case for CDS options, these option prices can be uniformly approximated as described in the next section.

The characteristics of the martingales M_t^S and M_t^X do not appear explicitly in the bond and CDS pricing formulas of Sections 2.2 and 2.3. This leaves the freedom to specify exogenous factors that feed into M_t^S and M_t^X . Such factors would be unspanned by the term structures of defaultable bonds and CDS and give rise to unspanned stochastic volatility, as described in (Filipović, Larsson, and Trolle 2016). Such stochastic volatility factors affect the distribution of future values of the survival and factor processes and therefore can be recovered from prices of CDS options and UCVA's. They also provide additional flexibility for fitting time series of bond prices and CDS spreads.

3 The Linear Hypercube Model

The linear hypercube (LHC) model assumes that the survival process S_t is absolutely continuous, as in Remark 2.1, and that the factor process X_t is

diffusive and takes values in a hypercube whose edges' length is given by S_t . More formally the state space of (S_t, X_t) is given by

$$E = \{(s, x) \in \mathbb{R}^{1+m} : s \in (0, 1] \text{ and } x \in [0, s]^m\}.$$

The dynamics of (S_t, X_t) is

$$\begin{aligned} dS_t &= -\gamma^\top X_t dt \\ dX_t &= (\beta S_t + B X_t) dt + \Sigma(S_t, X_t) dW_t \end{aligned} \tag{3}$$

for some $\gamma \in \mathbb{R}_+^m$ and some m -dimensional Brownian motion W_t , and where the dispersion matrix $\Sigma(S_t, X_t)$ is given by

$$\Sigma(s, x) = \text{diag} \left(\sigma_1 \sqrt{x_1(s - x_1)}, \dots, \sigma_m \sqrt{x_m(s - x_m)} \right) \tag{4}$$

with volatility parameters $\sigma_1, \dots, \sigma_m \geq 0$.

Let (S_t, X_t) be an E -valued solution of (3). It is readily verified that S_t is non-increasing and that the parameter γ controls the speed at which it decreases

$$0 \leq \gamma^\top X_t \leq \gamma^\top \mathbf{1} S_t$$

such that

$$0 \leq \lambda_t \leq \gamma^\top \mathbf{1} \text{ and } S_t \geq S_0 e^{-\gamma^\top \mathbf{1} t} > 0 \text{ for any } t \geq 0. \tag{5}$$

Note that the default intensity upper bound $\gamma^\top \mathbf{1}$ is a free parameter in this model. The following theorem gives conditions on the parameters such that the LHC model (3) is well defined.

Theorem 3.1. *Assume that, for all $i = 1, \dots, m$,*

$$\beta_i - \sum_{j \neq i} B_{ij}^- \geq 0, \tag{6}$$

$$\gamma_i + B_{ii} + \beta_i + \sum_{j \neq i} (\gamma_j + B_{ij})^+ \leq 0. \tag{7}$$

Then for any initial law of (S_0, X_0) with support in E there exists a unique in law E -valued solution (S_t, X_t) of (3). It satisfies the boundary non-attainment, for some $i = 1, \dots, m$,

(i) $X_{it} > 0$ for all $t \geq 0$ if $X_{i0} > 0$ and

$$\beta_i - \sum_{j \neq i} B_{ij}^- \geq \frac{\sigma_i^2}{2}, \quad (8)$$

(ii) $X_{it} < S_t$ for all $t \geq 0$ if $X_{i0} < S_0$ and

$$\gamma_i + B_{ii} + \beta_i + \sum_{j \neq i} (\gamma_j + B_{ij})^+ \leq -\frac{\sigma_i^2}{2}. \quad (9)$$

The state space E is a regular $(m + 1)$ -dimensional hyperpyramid. Figure 1 shows E when $m = 1$ and illustrates the drift inward pointing conditions (6)–(7) at the boundaries of E .

Remark 3.2. *One may consider a more generic specification of the process in which the dispersion matrix is given by*

$$\Sigma(s, x) = \text{diag} \left(\sigma_1 \sqrt{x_1(L_1 s - x_1)}, \dots, \sigma_m \sqrt{x_m(L_m s - x_m)} \right)$$

for some positive constants L_1, \dots, L_m . Lemma A.3 in the appendix shows that such model is observationally equivalent to the above specification.

In Appendix C we describe all possible market price of risk specifications under which the drift function of (S_t, X_t) remains linear.

3.1 One-Factor LHC Model

The default intensity of the one-factor LHC model, $m = 1$, has autonomous dynamics of the form

$$d\lambda_t = (\lambda_t^2 + B\lambda_t + \beta\gamma) dt + \sigma \sqrt{\lambda_t(\gamma - \lambda_t)} dW_t.$$

The diffusion function of λ_t is the same as the diffusion function of a Jacobi process taking values in the compact interval $[0, \gamma]$. However, the drift of λ_t includes a quadratic term that is neither present in Jacobi nor in affine processes.¹ Conditions (6)–(7) in Theorem 3.1 rewrite

$$\beta \geq 0 \quad \text{and} \quad (\gamma + B + \beta) \leq 0.$$

¹(Delbaen and Shirakawa 2002) construct a model in which the short rate is driven by a standard Jacobi process, the risk-free bond prices are then given by weighted series of Jacobi polynomials.

That is, the drift of λ_t is nonnegative at $\lambda_t = 0$ and nonpositive at $\lambda_t = \gamma$. We can factorize the drift as

$$\lambda_t^2 + B\lambda_t + \beta\gamma = (\lambda_t - \ell_1)(\lambda_t - \ell_2)$$

for some roots $0 \leq \ell_1 \leq \gamma \leq \ell_2$. This way, the default intensity mean-reverts to ℓ_1 . The corresponding original parameters are $B = -(\ell_1 + \ell_2)$ and $\beta\gamma = \ell_1\ell_2$, so that the drift of the factor X_t reads

$$\beta S_t + BX_t = (\ell_1 + \ell_2) \left(\frac{\ell_1\ell_2}{\gamma(\ell_1 + \ell_2)} S_t - X_t \right). \quad (10)$$

As a sanity check we verify that the constant default intensity case $\lambda_t = \gamma$ is nested as a special case. This is equivalent to have $X_t = S_t$ which can be obtained by specifying the dynamics $dX_t = -\gamma X_t dt$ for the factor process and the initial condition $X_0 = 1$. This corresponds to the roots $\ell_1 = 0$ and $\ell_2 = \gamma$.

The dynamics of the standard one-factor affine model on \mathbb{R}_+ is

$$d\lambda_t = \ell_2(\lambda_t - \ell_1)dt + \sigma\sqrt{\lambda_t}dW_t,$$

where ℓ_2 is the mean-reversion speed and ℓ_1 the mean-reversion level. Figure 2 shows the drift and diffusion functions of the default intensity for the one-factor LHC and affine model. The drift function is linear in the affine model whereas it is quadratic in the LHC model. However, for reasonable parameters values, the drift functions look similar when the default intensity is smaller than the mean-reversion level ℓ_1 . On the other hand, the force of mean-reversion above ℓ_1 is smaller and concave in the LHC model. The diffusion function is strictly increasing and concave for the affine model whereas it has a concave semi-ellipse shape in the LHC model. The diffusion functions have the same shape on $[0, \gamma/2]$ but typically do not scale equivalently in the parameter σ . However, the parameter γ can always be set sufficiently large so that the likelihood of λ_t going above $\gamma/2$ is arbitrarily small.

3.2 Option Approximation

We saw in Sections 2.3 and 2.4 that the pricing of a CDS option or a UCVA boils down to computing a \mathcal{F}_t -conditional expectation of the form

$$\Phi(f; t, T) = \mathbb{E}[f(S_T, X_T) | \mathcal{F}_t]$$

for some continuous function $f(s, x)$ on E . We now show how to approximate $\Phi(f; t, T)$ in analytical form by means of a polynomial approximation of $f(s, x)$ building on the results of (Filipović and Larsson 2016).

To this end, denote by $\text{Pol}_n(E)$ the set of polynomials $p(s, x)$ on E of degree n or less. It is readily seen that the generator of (S_t, X_t) ,

$$\mathcal{G}f(s, x) = (-\gamma^\top x, (\beta s + Bx)^\top) \nabla f(s, x) + \frac{1}{2} \sum_{i=1}^m \frac{\partial^2 f(s, x)}{\partial x_i^2} \sigma_i^2 x_i (s - x_i), \quad (11)$$

is polynomial in the sense that

$$\mathcal{G}\text{Pol}_n(E) \subset \text{Pol}_n(E) \quad \text{for any } n \in \mathbb{N}.$$

Let $N_n = \binom{n+1+m}{n}$ denote the dimension of $\text{Pol}_n(E)$ and fix a polynomial basis $\{h_1, \dots, h_{N_n}\}$ of $\text{Pol}_n(E)$. We define the function of (s, x)

$$H_n(s, x) := (h_1(s, x), \dots, h_{N_n}(s, x))^\top$$

with values in \mathbb{R}^{N_n} . There exists a unique matrix representation G_n of $\mathcal{G} \mid_{\text{Pol}_n(E)}$ with respect to this polynomial basis such that for any $p \in \text{Pol}_n(E)$ we can rewrite

$$\mathcal{G}p(s, x) = H_n(s, x)^\top G_n \vec{p}$$

where \vec{p} is the coordinate representation of p . (Filipović and Larsson 2016, Theorem (3.1)) then states that for any $t \leq T$ we have

$$\mathbb{E} [p(S_T, X_T) \mid \mathcal{F}_t] = H_n(S_t, X_t)^\top e^{G_n(T-t)} \vec{p}. \quad (12)$$

Remark 3.3. *The choice for the basis $H_n(x, s)$ of $\text{Pol}_n(E)$ is arbitrary and one may simply consider the monomial basis,*

$$H_n(s, x) = \{1, s, x_1, \dots, x_m, s^2, sx_1, x_1^2, \dots, x_m^n\}$$

in which G_n is block-diagonal. There are efficient algorithms to compute the matrix exponential $e^{G_n(T-t)}$, see for example (Higham 2008). Note that only the action of the matrix exponential is required, that is $e^{G_n(T-t)} \vec{p}$ for some $p \in \text{Pol}_n(E)$, for which specific algorithms exist as well, see for examples (Al-Mohy and Higham 2011) and (Sidje 1998) and references within.

Now let $\epsilon > 0$. From the Stone-Weierstrass approximation theorem (Rudin 1974, Theorem 5.8) there exists a polynomial $p \in \text{Pol}_n(E)$ for some n such that

$$\sup_{(s,x) \in E} |f(s,x) - p(s,x)| \leq \epsilon. \quad (13)$$

Combining (12) and (13) we obtain the desired approximation of $\Phi(f; t, T)$.

Theorem 3.4. *Let $p \in \text{Pol}_n(E)$ be as in (13). Then $\Phi(f; t, T)$ is uniformly approximated by*

$$\sup_{t \leq T} \|\Phi(f; t, T) - H_n(S_t, X_t)^\top e^{G_n(T-t)} \vec{p}\|_{L^\infty} \leq \epsilon. \quad (14)$$

The approximating polynomial p in (13) needs to be found case by case. We illustrate this for the CDS option (2). It is readily seen that the support of the random variable $Z(T_0, T_N, k) = \eta^\top \begin{pmatrix} S_{T_0} \\ X_{T_0} \end{pmatrix}$ is contained in $[a, b]$ where we write $\eta = \gamma_{\text{prot}}(T_0, T_0, T_N) - k \gamma_{\text{prem}}(T_0, T_0, T_N)$ and

$$a = \mathbf{1}^\top \eta^- \quad \text{and} \quad b = \mathbf{1}^\top \eta^+$$

where $\mathbf{1}$ is the $(1 + m)$ -dimensional vector with coordinates equal to one. So we need to approximate the real function z^+ on $[a, b]$ by a polynomial $q \in \text{Pol}_n([a, b])$, which can be done by Chebyshev interpolation as outlined in Appendix B. The approximating polynomial $p \in \text{Pol}_n(E)$ is then given by $p(s, x) = q(\eta^\top(s, x))$.

Remark 3.5. *Approximating the payoff function $f(s, x)$ on a strict subset of the state space E is sufficient to approximate an option price. Indeed, for any times $t \leq u \leq T$ the process (S_u, X_u) takes values in*

$$\{(s, x) \in E : S_t \geq s \geq e^{-\gamma^\top \mathbf{1}(T-t)} S_t\} \subset E.$$

A polynomial approximation on a compact set smaller than E can be expected to be more precise and, as a result, to produce a more accurate price approximation.

4 Empirical Analysis

We show that LHC models are able to capture complex term structure dynamics and that options prices can be accurately approximated. First, we

fit parsimonious multi-factor models to CDS data and discuss the estimated parameters and factors. Second, we accurately approximate the price of CDS options at different moneyness for the three-factor model. Finally, we estimate corresponding affine default intensity models and compare the fits.

4.1 Data

The empirical analysis is based on composite CDS spread data from Markit which are essentially averaged quotes provided by major market makers. At each date we include the available spreads on contracts with maturities of 1, 2, 3, 4, 5, 7, and 10 years. We choose the CDS spreads on a major financial player, namely *JP Morgan Chase & Co*, with full restructuring clause (CR) meaning that any restructuring event qualifies as credit event. The data set contains 565 weekly observations summing up to 3897 observed CDS spreads. The sample starts on July 20, 2004 and ends on May 11, 2015. The choice of this time series is motivated by the important fluctuations in the spreads and their volatility, and the prominence of the underlying entity which we suppose ensures a sufficient trading volume and thus meaningful quotes.

Time series of the 1-year, 5-year, and 10-year CDS spreads are displayed in Figure 3, as well as the relative changes on the 5-year versus 1-year CDS spread. The time series can be split in three time periods. The first period, before the subprime crisis, exhibits low spreads in contango and low volatility. The second period, during the subprime crisis, exhibits high volatility with skyrocketing spreads temporarily in backwardation. The more recent period is characterized by a steep contango and a lot of volatility. Figure 3 also shows that CDS spread changes are strongly correlated across maturities.

The summary statistics are reported in Table 1. A principal component analysis reveals that close to 95% of the variation in the change of the term structure can be explained by the first factor, the second and third factors respectively explaining 3% and 1%. This suggests that a two- or three-factor model should be able to capture most of the term structure dynamics.

4.2 Model Specifications

We focus on the LHC model of Section 3 with two and three factors. We set $\gamma = \gamma_1 \mathbf{e}_1$, for some $\gamma_1 \geq 0$, and consider a cascading structure of the form

$$dX_{it} = \kappa_i(\theta_i X_{(i+1)t} - X_{it}) dt + \sigma_i \sqrt{X_{it}(S_t - X_{it})} dW_{it} \quad (15)$$

for $i = 1, \dots, m - 1$ and

$$dX_{mt} = \kappa_m(\theta_m S_t - X_{mt}) dt + \sigma_m \sqrt{X_{mt}(S_t - X_{mt})} dW_{mt} \quad (16)$$

for some parameters $\kappa, \theta, \sigma \in \mathbb{R}_+^m$ satisfying

$$\theta_i \leq 1 - \frac{\gamma_1}{\kappa_i} \quad \text{for } i = 1, \dots, m.$$

It is straightforward to verify that conditions (6) and (7) hold.

This specification allows for default intensity values to be persistently close to zero over extended periods of time. It also allows to work with a multidimensional model parsimoniously as the number of free parameters is equal to $3m+1$ whereas it is equal to $3m+m^2$ for the generic LHC model. This structure also appeared naturally in repeated models estimations on multiple CDS term structures. Note that the default intensity is then proportional to the first factor and given by $\lambda_t = \gamma_1 X_{1t}/S_t$.

We denote the two-factor and three-factor linear hypercube cascade models by LHCC(2) and LHCC(3), respectively. In addition, we estimate a three-factor model with parameter fixed $\gamma_1 = 25\%$ that we denote by LHCC(3)*. This parameter value is about twice as large as the estimated γ_1 from the LHCC(3) model. We estimate the constrained model in order to determine whether the choice of the default intensity upper bound is determinant for the empirical results.

We set the market price of risk for the Brownian motions to zero. Numerous estimations have shown that the inclusion of non-zero market price of risk parameters does not improve the fit for the LHCC models. Market price of risk specifications that have been explored are described in Appendix C.

We set the risk-free rate equal to the average 5-year risk-free yield over the sample, $r = 2.53\%$. We make the usual assumption that the recovery rate is equal to $\delta = 40\%$.

4.3 Maximum Likelihood Estimates

Quasi-maximum likelihood parameters estimates and their asymptotic standard deviations are reported in Table 3. The one-week joint conditional density of the survival process and the factors is approximated with a Gaussian density and filtered from observed CDS spreads using the unscented Kalman filter (UKF) at each date in the sample. The UKF has been shown

to perform better than alternative Kalman filters when approximating dynamic systems with non-linear dependence in the states, see (Christoffersen, Dorion, Jacobs, and Karoui 2014). Recall that the observable CDS spreads are linear-rational in (S_t, X_t) . We assume that the spreads are observed with i.i.d. normal errors having a variance equal to the parameter σ_{cds}^2 which is also estimated. The standard errors have been computed using all possible lagged autocorrelation, and are therefore more conservative than the Newey-West standard errors.

The parameters estimates take similar values across models, and the majority of them are statistically significant. The parameter γ_1 is equal to 7.44% and 11.14% for the unconstrained two- and tree-factor model, respectively. These bounds are arguably larger than the average level of default intensity extracted from CDS spreads which is around 0.6%. They are also larger than the maximal default intensity values attained during the financial crisis which topped slightly below 6%. The speed of mean-reversion is always smaller for the m -th factor. This suggests that the long run mean-reversion level of the default intensity may vary over prolonged periods of time. The volatility parameters take large absolute values which is not surprising given that they scale the diffusion matrix at a smaller rate than affine models, see Figure 2.

The CDS spreads having a closed-form expression, it takes less than a second to filter the entire time series of states on a single CPU core with a standard computer. This is noticeable as the UKF is a time consuming methodology, requiring at each date to compute CDS spreads at many fictive states positions. As a result, the estimation procedure is swift.

4.4 Specification Analysis

The fitted factors extracted from the quasi-maximum likelihood estimations are used as input to compute the fitted spreads. With the fitted spreads we compute the fit errors for each date and maturity. Not surprisingly the more flexible specification LHCC(3) seems to perform the best. Estimating the default intensity upper bound γ_1 instead of setting an arbitrarily large value improves the calibration. Table 2 reports summary statistics of the errors by maturity. The LHCC(3) model has the smaller RMSE for each maturity. In particular, its overall RMSE is half the one of the two-factor model. The LHCC(3)* model faces difficulties in reproducing long-term spreads as, for example, its average error is twice larger than the one of the unconstrained

LHCC(3) for the 10-year maturity spread. Figure 4 displays the fitted spreads and the RMSE time series. Again, the LHCC(3) appears to have the smallest level of errors over time. The two alternative models do not perform as good during the low spreads period before the financial crisis, and during the recent volatile period. Note that the quasi-maximum likelihood estimate of the observation errors σ_{cds} is also smaller for the unconstrained three-factor model.

4.5 Factors

Figure 5 shows the estimated factors. They are remarkably similar across the different specifications. The default intensity explodes and the survival process decreases rapidly during the financial crisis and, to a smaller extent, during the European sovereign debt crisis. The m -th factor controls the long term default intensity level. The second factors controls the medium term behavior of the term-structure of credit risk in the three-factor models.

The two-factor model requires an almost equal to zero default intensity to capture the steep contango of the term structure at the end of the period, even lower than before the financial crisis. This seems counterfactual and illustrates the limitations of the two-factor model to capture a changing dynamics.

The m -th factor visits the second half of its support $[0, S_t]$ and appears to stabilize in this region for the three models. Having a lower volatility at both ends of this support, the process X_{mt} tends to be pushed away from the center. As a result, some factors alternate visits to the lower part and upper part of their supports, which may facilitate the shift and persistence between different shapes of term structure.

4.6 Pricing CDS Options

We explore the price approximation of CDS options at different moneyness using the methodology presented in Section 3.2 for the estimated LHCC(3) model. We consider three strike spreads equal to 25, 50, and 100 basis points. On the observation date, January 5th 2015, the CDS spreads in the sample took values ranging from 25 to 109 basis points.

The CDS option payoff function is interpolated by means of Chebyshev polynomials. Details are given in Appendix B. The first row in Figure 6 shows the payoff approximation $q(z)$ for the polynomial orders $N \in \{2, 10, 20\}$ and

different strike spreads. We observe that a more accurate approximation of the hockey stick payoff function is obtained by increasing the order N , especially around the kink. The width of $Z(T_0, T_N, k)$'s support increases with the strike spread k . Hence, the uniform error bound should be expected to be larger for out of the money options. This is confirmed by the second row of Figure 6 that shows the error bound as a function of the approximation order N . We remind that this error bound is typically non tight.

The third row of Figure 6 shows the price approximation as a function of the polynomial order, up to $N = 35$. The price approximations oscillate and then stabilize at the true option prices. The amplitudes of these fluctuations on the range of price approximation orders $N \in [20, 35]$ are small with an average of 0.10, 1.06, and 0.68 basis points around the range price approximation mean for the strikes 25, 50, and 100 basis points respectively.

Finally, note that the dimension of the polynomial basis becomes a computational problem when both the expansion order N and the number of factors $1 + m$ are large. For example, for $N = 35$ and $m = 3$ the basis H_n has dimension 82'251 whereas it has only dimension 662 for $m = 2$.

4.7 Comparison with Affine Default Intensity Models

We fit affine default intensity models to the same sample of CDS spreads. For the same reasons as in Section 4.2 and for the sake of comparison, we specify a cascading structure and set the market price of risk to zero. The default intensity is defined to be $\lambda_t = X_{1t}$ where the process $X_t \in \mathbb{R}_+^m$ is given by

$$dX_{it} = \kappa_i(\theta_i + X_{(i+1)t} - X_{it}) dt + \sigma_i \sqrt{X_{it}} dW_{it}$$

for $i = 1, \dots, m - 1$ and

$$dX_{mt} = \kappa_m(\theta_m - X_{mt}) dt + \sigma_m \sqrt{X_{mt}} dW_{mt}$$

for some parameters $\kappa, \theta, \sigma \in \mathbb{R}_+^m$. We label the models with two and three factors AC(2) and AC(3) respectively. Note that the parameters are left unconstrained for the affine models, beyond nonnegativity, and that the LHCC models have an additional parameter γ_1 fixing the default intensity upper bound.

Table 4 reports summary statistic of the fitted errors and shows that the LHCC and AC models perform equivalently. Yet, the affine models seem to perform marginally better on average. However, an important number of

parameter estimates appear to be weakly identified for the affine models as can be seen in Table 5. This suggests that either the specification can be refined or that it is difficult to identify. The estimated parameters convey that the process X_t is mean-reverting at an extremely low speed. The estimated factors displayed in Figure 7 show that they have a lot of similarities with those extracted with the LHCC models.

The estimation procedure is more than ten times slower with affine default intensity models than with linear credit risk models. Indeed, the different contingent bond prices necessary to compute the CDS spreads have to be computed numerically for the AC models, as discussed in Section 2.2. In this application the nonlinear ODEs are numerically solved only once for each parameter set by means of a Runge-Kutta 4th order integration.

5 Conclusion

We introduce the class of linear credit risk models in which the background survival probability of a firm and its factors have a linear drift. The prices of defaultable bonds and CDS become linear in the factors. We define the linear hypercube (LHC) model with diffusive factor process that has quadratic diffusion function and takes values in a compact state space. These features are employed to develop an efficient European option pricing methodology. Empirical analysis shows that the LHC model is able to capture the complex CDS term structure dynamics. The prices of CDS options at different moneyness are accurately approximated.

Future research can follow different directions. Models with jumps or different state space specifications are to be explored. Work in progress includes the development of multi-firm models with correlated defaults. One particular direction is the construction of linear credit risk models in which survival processes are driven by common factors.

A Proofs

Proof of (1)

This follows as in (Filipović, Larsson, and Trolle 2016, Lemma 3).

Proof of Proposition 2.2

Proposition 2.2 is an immediate consequence of (1) and the following lemma.

Lemma A.1. *Let Y be a nonnegative \mathcal{F}_∞ -measurable random variable. For any $t \leq T < \infty$,*

$$\mathbb{E} [\mathbb{1}_{\{\tau > T\}} Y \mid \mathcal{G}_t] = \mathbb{1}_{\{\tau > t\}} \frac{1}{S_t} \mathbb{E} [S_T Y \mid \mathcal{F}_t].$$

Note that $T < \infty$ is essential unless we assume that $S_\infty = 0$.

Lemma A.1 follows from (Bielecki and Rutkowski 2002, Corollary 5.1.1). For the convenience of the reader we provide here a sketch of its proof. As in (Bielecki and Rutkowski 2002, Lemma 5.1.2) one can show that, for any nonnegative random variable Z , we have

$$\mathbb{E} [\mathbb{1}_{\{\tau > t\}} Z \mid \mathcal{H}_t \vee \mathcal{F}_t] = \mathbb{1}_{\{\tau > t\}} \frac{1}{S_t} \mathbb{E} [\mathbb{1}_{\{\tau > t\}} Z \mid \mathcal{F}_t].$$

Setting $Z = \mathbb{1}_{\{\tau > T\}} Y$ we can now derive

$$\begin{aligned} \mathbb{E} [\mathbb{1}_{\{\tau > T\}} Y \mid \mathcal{G}_t] &= \mathbb{E} [\mathbb{1}_{\{\tau > t\}} Y \mathbb{1}_{\{\tau > T\}} \mid \mathcal{G}_t] = \mathbb{1}_{\{\tau > t\}} \frac{1}{S_t} \mathbb{E} [\mathbb{1}_{\{\tau > T\}} Y \mid \mathcal{F}_t] \\ &= \mathbb{1}_{\{\tau > t\}} \frac{1}{S_t} \mathbb{E} [\mathbb{E} [\mathbb{1}_{\{\tau > T\}} \mid \mathcal{F}_\infty] Y \mid \mathcal{F}_t] = \mathbb{1}_{\{\tau > t\}} \frac{1}{S_t} \mathbb{E} [S_T Y \mid \mathcal{F}_t]. \end{aligned}$$

Proof of Proposition 2.3

The subsequent proofs build on the following lemma that follows from (Bielecki and Rutkowski 2002, Proposition 5.1.1).

Lemma A.2. *Let Z_t be a bounded \mathcal{F}_t -predictable process. For any $t \leq T < \infty$,*

$$\mathbb{E} [\mathbb{1}_{\{t < \tau \leq T\}} Z_\tau \mid \mathcal{G}_t] = \mathbb{1}_{\{t < \tau\}} \frac{1}{S_t} \mathbb{E} \left[\int_{(t, T]} -Z_u dS_u \mid \mathcal{F}_t \right].$$

Note that $T < \infty$ is essential unless we assume that $S_\infty = 0$.

We can now proceed to the proof of Proposition 2.3. The value of the contingent cash flow is given by the expression

$$C^D(t, T) = \mathbb{E} \left[e^{-r(\tau-t)} \mathbb{1}_{\{t \leq \tau \leq T\}} \mid \mathcal{G}_t \right]$$

By applying Lemma A.2 we get

$$\begin{aligned} C^D(t, T) &= \frac{\mathbb{1}_{\{\tau > t\}}}{S_t} \mathbb{E} \left[\int_t^T -e^{-r(s-t)} dS_s \mid \mathcal{F}_t \right] \\ &= \frac{\mathbb{1}_{\{\tau > t\}}}{S_t} \int_t^T e^{-r(s-t)} \mathbb{E} [\gamma^\top X_s \mid \mathcal{F}_t] ds \\ &= \frac{\mathbb{1}_{\{\tau > t\}}}{S_t} \int_t^T e^{-r(s-t)} \begin{pmatrix} 0 & \gamma^\top \end{pmatrix} e^{A(s-t)} \begin{pmatrix} S_t \\ X_t \end{pmatrix} ds \end{aligned}$$

where the second equality comes from the fact that $\int_0^t e^{-ru} dM_u^S$ is a martingale. The third equality follows from (1). The proposition now follows by observing that for any matrix A and real r we have $e^r e^A = e^{\text{diag}(r)+A}$, and that the matrix exponential integration can be computed explicitly as follows

$$\begin{aligned} \int_0^u e^{As} ds &= \int_0^u (I + As + A^2 \frac{s^2}{2} + \dots) ds = Iu + A \frac{u^2}{2} + A^2 \frac{u^3}{6} + \dots \\ &= A^{-1} (e^{Au} - I). \end{aligned}$$

Proof of Lemma 2.4

The value of this contingent bond is given by

$$C^{D\tau}(t, T) = \mathbb{E} [\tau e^{-r(\tau-t)} \mathbb{1}_{\{t < \tau \leq T\}} \mid \mathcal{G}_t]$$

which, by following the same step as in the previous proposition, is given by

$$C^{D\tau}(t, T) = \frac{\mathbb{1}_{\{\tau > t\}}}{S_t} \begin{pmatrix} 0 & \gamma^\top \end{pmatrix} \left(\int_t^T s e^{A_*(s-t)} ds \right) \begin{pmatrix} S_t \\ X_t \end{pmatrix}.$$

By change of variable $u = s - t$ the integral within parenthesis rewrites

$$\int_t^T s e^{A_*(s-t)} ds = \int_0^{T-t} u e^{A_* u} du + t \int_0^{T-t} e^{A_* u} du,$$

where the second term on the RHS is given in Lemma 2.3. The first term can be derived using integration by parts

$$\int_0^{T-t} u e^{A_* u} du = (T-t) A_*^{-1} e^{A_*(T-t)} - A_*^{-1} A_*^{-1} (e^{A_*(T-t)} - I).$$

Proof of Proposition 2.5

The calculations of the protection leg V_{prot} and the coupon part V_{coup} are straightforward. The accrued interest V_{ai} is given by the sum of contingent cash flows and of weighted zero-recovery coupon bonds. The series of contingent cash flow is in fact equal to a single contingent payment paying τ at default,

$$C^{D\tau}(t, T) = \sum_{k=1}^N \mathbb{E} \left[\tau e^{-r(\tau-t)} \mathbb{1}_{\{T_{k-1} < \tau \leq T_k\}} \mid \mathcal{G}_t \right] = \mathbb{E} \left[\tau e^{-r(\tau-t)} \mathbb{1}_{\{t < \tau \leq T\}} \mid \mathcal{G}_t \right].$$

Using the identity $\mathbb{1}_{\{T_{k-1} < \tau \leq T_k\}} = \mathbb{1}_{\{\tau > T_{k-1}\}} - \mathbb{1}_{\{\tau > T_k\}}$ we obtain that the second term of V_{ai} is given by

$$\begin{aligned} -\mathbb{E} \left[\sum_{k=1}^N e^{-r(\tau-t)} T_{k-1} \mathbb{1}_{\{T_{k-1} < \tau \leq T_k\}} \mid \mathcal{G}_t \right] &= \sum_{k=1}^N T_{k-1} (C^D(t, T_k) - C^D(t, T_{k-1})) \\ &= T_{N-1} C^D(t, T_N) - T_0 C^D(t, T_0) - \sum_{k=1}^{N-1} (T_k - T_{k-1}) C^D(t, T_k). \end{aligned}$$

Proof of Theorem 3.1

We define the bounded continuous map $(\mathcal{S}, \mathcal{X}) : R^{1+m} \rightarrow R^{1+m}$ by

$$\mathcal{S}(s, x) = s^+ \wedge 1, \quad \mathcal{X}_i(s, x) = x_i^+ \wedge s^+ \wedge 1, \quad i = 1, \dots, m,$$

such that $(\mathcal{S}, \mathcal{X})(s, x) = (s, x)$ on E . In a similar vein, extend the dispersion matrix $\Sigma(s, x)$ to a bounded continuous mapping $\Sigma((\mathcal{S}, \mathcal{X})(s, x))$ on \mathbb{R}^{1+m} . The stochastic differential equation (3) then extends to \mathbb{R}^{1+m} by

$$\begin{aligned} dS_t &= -\gamma^\top \mathcal{X}(S_t, X_t) dt \\ dX_t &= (\beta \mathcal{S}(S_t) + B \mathcal{X}(S_t, X_t)) dt + \Sigma((\mathcal{S}, \mathcal{X})(S_t, X_t)) dW_t. \end{aligned} \tag{17}$$

Since drift and dispersion of (17) are bounded and continuous on \mathbb{R}^{1+m} , there exists a weak solution (S_t, X_t) of (17) for any initial law of (S_0, X_0) with support in E , see (Karatzas and Shreve 1991, Theorem V.4.22).

We now show that any weak solution (S_t, X_t) of (17) with $(S_0, X_0) \in E$ stays in E ,

$$(S_t, X_t) \in E \text{ for all } t \geq 0. \tag{18}$$

To this end, for $i = 1, \dots, m$, note that

$$\Sigma_{ii}((\mathcal{S}, \mathcal{X})(s, x)) = 0 \text{ for all } (s, x) \text{ with } x_i \leq 0 \text{ or } x_i \geq s. \quad (19)$$

Conditon (6) implies that

$$(\beta\mathcal{S}(s) + B\mathcal{X}(s, x))_i \geq 0 \text{ for all } (s, x) \text{ with } x_i \leq 0. \quad (20)$$

For $\delta, \epsilon > 0$ we define

$$\tau_{\delta, \epsilon} = \inf \{t \geq 0 \mid X_{it} \leq -\epsilon \text{ and } X_{is} < 0 \text{ for all } s \in [t - \delta, t]\}.$$

Then on $\{\tau_{\delta, \epsilon} < \infty\}$ we have, in view of (19) and (20),

$$0 > X_{i\tau_{\delta, \epsilon}} - X_{i\tau_{\delta, \epsilon} - \delta} = \int_{\tau_{\delta, \epsilon} - \delta}^{\tau_{\delta, \epsilon}} (\beta\mathcal{S}(S_u) + B\mathcal{X}(S_u, X_u))_i du \geq 0,$$

which is absurd. Hence $\tau_{\delta, \epsilon} = \infty$ a.s. and therefore $X_{it} \geq 0$ for all $t \geq 0$. Similarly, conditon (7) implies that

$$-\gamma^\top \mathcal{X}(s, x) - (\beta\mathcal{S}(s) + B\mathcal{X}(s, x))_i \geq 0 \text{ for all } (s, x) \text{ with } x_i \geq s. \quad (21)$$

Using the same argument as above for $S_t - X_{it}$ in lieu of X_{it} , and (21) in lieu of (20), we see that $S_t - X_{it} \geq 0$ for all $t \geq 0$. Finally, note that $0 \leq \gamma^\top \mathcal{X}(s, x) \leq \gamma^\top \mathbf{1}s^+$ for all (s, x) , and thus $1 \geq S_t \geq e^{-\gamma^\top \mathbf{1}t} > 0$ for all $t \geq 0$. This proves (18) and thus the existence of an E -valued solution of (3).

Uniqueness in law of the E -valued solution (S_t, X_t) of (3) follows from (Filipović and Larsson 2016, Theorem 4.2) and the fact that E is relatively compact.

The boundary non-attainment conditions (8)–(9) follow from (Filipović and Larsson 2016, Theorem 5.7(i) and (ii)) for the polynomials $p(s, x) = x_i$ and $s - x_i$, for $i = 1, \dots, m$.

Proof of Remark 3.2

The claim in Remark 3.2 follows from the following lemma.

Lemma A.3. *The process (S_t, X_t) with drift as in (3) and dispersion matrix given by $\text{diag} \left(\sigma_1 \sqrt{x_1(L_1 s - x_1)}, \dots, \sigma_m \sqrt{x_m(L_m s - x_m)} \right)$ for some positive constants L_1, \dots, L_m is observationally equivalent to the one with $L_i = 1$ for all $i = 1, \dots, m$.*

Proof. This directly follows by applying the change of variable $X'_{it} = X_{it}/L_i$ for each $i = 1, \dots, m$. The dynamics of (S_t, X'_t) then rewrites

$$\begin{aligned} dS_t &= -\gamma'^\top X'_t dt \\ dX'_t &= (\beta' S_t + B' X'_t) dt + \Sigma(S_t, X'_t) dW_t \end{aligned}$$

with

$$\begin{aligned} \gamma' &= \text{diag}(L_1, \dots, L_m) \gamma, \\ \beta' &= \text{diag}(L_1, \dots, L_m) \beta, \\ B' &= \text{diag}(1/L_1, \dots, 1/L_m) B \text{diag}(L_1, \dots, L_m), \end{aligned}$$

and where the diffusion matrix is given by Equation (4). \square

B Chebyshev Interpolation

This section describes how to perform a Chebyshev interpolation of an arbitrary function on an interval $[a, b] \subset \mathbb{R}$. The Chebyshev polynomials of the first kind take values in $[-1, 1]$ but can be shifted and scaled so as to form a basis of $[a, b]$. In this case they are given by the following recursion formula,

$$\begin{aligned} T_0^{a,b}(x) &= 1 \\ T_1^{a,b}(x) &= \frac{x - \mu}{\sigma} \\ T_{n+1}^{a,b}(x) &= \frac{2(x - \mu)}{\sigma} T_n^{a,b}(x) - T_{n-1}^{a,b}(x) \end{aligned}$$

with $\mu = (a + b)/2$ and $\sigma = (b - a)/2$. The Chebyshev interpolation of the function f is the unique polynomial $p_N(x_j)$ of order N that satisfies

$$p_N(x_j) = f(x_j) \quad \text{for any } j = 0, \dots, N$$

where the Chebyshev node x_j is given by

$$x_j = \mu + \sigma \cos \left(\frac{(1/2 + j)\pi}{N + 1} \right).$$

The interpolation can also be written directly

$$p_N(x) = \sum_{n=0}^N c_n T_n^{a,b}(x)$$

where the coefficients c_n are explicit

$$c_n = \frac{2}{N+1} \sum_{j=0}^N f(x_j) T_n^{a,b}(x_j).$$

The coefficients can be computed in an effective way by applying Clenshaw's method, or by applying discrete cosine transform. This straightforward interpolation has the advantage to avoid the Runge's phenomena, which means that if f is $N+1$ differentiable we have

$$\|f - p_N\|_{L^\infty} \leq \frac{1}{2^N(N+1)!} \|f\|_{L^\infty}$$

where $\|\cdot\|_{L^\infty}$ denotes the super norm over $[a, b]$, that is for any function g

$$\|g\|_{L^\infty} = \sup_{x \in [a, b]} |g(x)|.$$

Unfortunately in our case the payoff function $(x)^+$ is not differentiable at zero so that we report an error bound ϵ computed numerically. The Chesyshev interpolation can similarly be applied to multidimensional functions. For example, (Gaß, Glau, Mahlstedt, and Mair 2015) applies multidimensional Chebyshev interpolation to interpolate over the parameters space an objective function. Note that alternative polynomial expansion, such as Generalized Fourier Series with orthonormal polynomial basis, may provide a smaller error bound but may also come at a higher computational cost.

C Market Price of Risk Specifications

We discuss market price of risk (MPR) specifications such that X_t has a linear drift also under the real-world measure $\mathbb{P} \sim \mathbb{Q}$. This may further facilitate the empirical estimation of the LHC model.

Let $\Lambda(S_t, X_t)$ denote the MPR such that the drift of X_t under \mathbb{P} becomes

$$\mu_t^{\mathbb{P}} = \beta S_t + B X_t + \Sigma(S_t, X_t) \Lambda(S_t, X_t).$$

It is linear in (S_t, X_t) of the form

$$\mu_t^{\mathbb{P}} = \beta^{\mathbb{P}} S_t + B^{\mathbb{P}} X_t,$$

for some vector $\beta^{\mathbb{P}} \in \mathbb{R}^m$ and matrix $B^{\mathbb{P}} \in \mathbb{R}^{m \times m}$, if and only if

$$\Lambda_i(s, x) = \frac{((\beta^{\mathbb{P}} - \beta)s + (B^{\mathbb{P}} - B)x)_i}{\sigma_i \sqrt{x_i(s - x_i)}}, \quad i = 1, \dots, m. \quad (22)$$

In order that $\Lambda(S_t, X_t)$ is well defined and induces an equivalent measure change, that is, the candidate Radon–Nikodym density process

$$\exp \left(\int_0^t \Lambda(S_u, X_u) dW_u - \frac{1}{2} \int_0^t \|\Lambda(S_u, X_u)\|^2 du \right) \quad (23)$$

is a uniformly integrable \mathbb{Q} -martingale, we need that (S_t, X_t) does not attain all parts of the boundary of E . This is clarified by the following theorem, which follows from (Cheridito, Filipović, and Yor 2005).

Theorem C.1. *The MPR $\Lambda(S_t, X_t)$ in (22) is well defined and induces an equivalent measure $\mathbb{P} \sim \mathbb{Q}$ with Radon–Nikodym density process (23) if, for all $i = 1, \dots, m$, $X_{i0} \in (0, S_0)$ and (8)–(9) hold for the \mathbb{Q} -drift parameters B, β and for the \mathbb{P} -drift parameters $B^{\mathbb{P}}, \beta^{\mathbb{P}}$ in lieu of B, β .*

If, for some $i = 1, \dots, m$, $B_{ij}^{\mathbb{P}} = B_{ij}$ for all $j \neq i$ and

(i) $\beta_i^{\mathbb{P}} = \beta_i$, *such that*

$$\Lambda_i(s, x) = \frac{(B_{ii}^{\mathbb{P}} - B_{ii})\sqrt{x_i}}{\sigma_i \sqrt{s - x_i}},$$

then it is enough if $X_{i0} \in [0, S_0)$ instead of $X_{i0} \in (0, S_0)$ and (6) instead of (8) holds for B_{ij}, β_i , and thus for $B_{ij}^{\mathbb{P}}, \beta_i^{\mathbb{P}}$.

(ii) $\beta_i^{\mathbb{P}} - \beta_i = B_{ii}^{\mathbb{P}} - B_{ii}$, *such that*

$$\Lambda_i(s, x) = \frac{(B_{ii}^{\mathbb{P}} - B_{ii})\sqrt{s - x_i}}{\sigma_i \sqrt{x_i}},$$

then it is enough if $X_{i0} \in (0, S_0]$ instead of $X_{i0} \in (0, S_0)$ and (7) instead of (9) holds for B_{ij}, β_i , and thus for $B_{ij}^{\mathbb{P}}, \beta_i^{\mathbb{P}}$.

The assumption of linear-drift preserving change of measure is often made for parsimony and to facilitate the empirical estimation procedure. For example, the specification of MPRs that preserve the affine nature of risk-factors has been theoretically and empirically investigated in (Duffee 2002), (Duarte 2004), and (Cheridito, Filipović, and Kimmel 2007) among others.

	1 yr	2 yrs	3 yrs	4 yrs	5 yrs	7 yrs	10 yrs
Mean	37.96	46.55	55.84	69.45	74.57	83.08	90.95
Vol	40.49	40.01	40.64	41.77	43.58	43.01	42.67
Min	2.28	3.61	6.64	9.45	11.56	15.36	20.74
Max	294.14	257.52	248.50	243.24	250.23	242.73	235.90

Table 1: Summary statistics for the CDS spreads on JPMorgan Chase & Co. The sample contains 565 weekly observations starting on July 20, 2004 and ending on May 11, 2015 summing to 3897 CDS spreads. The CDS term structure experienced important variations in shape and level over this period of time.

Specification		all	1 yr	2 yrs	3 yrs	4 yrs	5 yrs	7 yrs	10 yrs
LHCC(2)	RMSE	5.08	4.30	4.59	5.36	6.19	5.98	2.67	5.71
	Vol	5.01	4.21	3.43	4.61	6.11	5.59	2.54	5.60
	Min	-19.90	-17.51	-14.99	-15.67	-15.11	-10.63	-13.28	-19.90
	Max	27.51	27.51	9.84	12.94	15.14	17.70	11.95	11.14
	Med	-0.46	-0.36	-2.40	-1.20	0.08	1.71	0.31	-0.76
LHCC(3)	RMSE	2.53	1.93	2.56	2.36	2.70	3.65	2.21	1.86
	Vol	2.51	1.75	2.30	1.77	2.50	3.31	2.19	1.82
	Min	-10.23	-8.21	-9.31	-8.29	-9.31	-9.19	-10.23	-7.16
	Max	14.45	14.45	8.48	4.38	3.96	11.02	7.82	6.80
	Med	-0.18	0.61	-0.42	-1.47	-0.28	1.45	-0.85	-0.03
LHCC(3)*	RMSE	3.77	2.48	2.25	3.59	5.03	4.77	2.43	4.73
	Vol	3.74	2.17	1.88	2.89	4.67	4.67	2.41	4.74
	Min	-16.26	-8.92	-12.83	-9.63	-16.26	-14.23	-10.91	-14.01
	Max	21.13	9.37	8.79	7.83	3.91	11.24	18.25	21.13
	Med	-0.25	0.86	-0.87	-1.47	0.21	2.18	-0.64	-0.84

Table 2: Comparison of fitted spread errors for the LHC models.

The table reports the minimal, maximal, median, and root mean squared errors by maturity over the entire time period for the three different specifications. The three-factor models perform better for all maturities than the two-factor model, and estimating the default intensity upper bound leads to a more accurate fit.

	LHCC(2)	LHCC(3)	LHCC(3)*
γ_1	0.0744 (0.0023)	0.1114 (0.0005)	0.2500
κ_1	0.2231 (0.0215)	0.4233 (0.0099)	0.4910 (0.0097)
κ_2	0.0974 (0.0078)	0.3283 (0.0061)	0.4729 (0.0105)
κ_3		0.1121 (0.0047)	0.3333 (0.0247)
θ_1	0.6664 (0.0159)	0.7368 (0.0033)	0.4908 (0.0038)
θ_2	0.2360 (0.0067)	0.6606 (0.0025)	0.4714 (0.0033)
θ_3		0.0054 (0.0121)	0.2498 (0.0068)
σ_1	1.1331 (0.1150)	0.4499 (0.0450)	0.3345 (0.0695)
σ_2	0.7090 (0.0734)	1.0084 (0.1052)	0.7928 (0.1034)
σ_3		0.5224 (0.0439)	1.0689 (0.1414)
X_{10}	0.2328 (0.6681)	0.0031 (0.0023)	0.0101 (0.0157)
X_{20}	0.5207 (0.3165)	0.0774 (0.0429)	0.0300 (0.0263)
X_{30}		0.1055 (0.0422)	0.1500 (0.1086)
σ_{cds}	5.2783 (0.0000)	2.8082 (0.0000)	4.0126 (0.0000)
r	0.0253	0.0253	0.0253
$\log\mathcal{L}$	22,901	24,809	23,728
AIC	-45,782	-49,589	-47,430
BIC	-45,720	-49,502	-47,349

Table 3: LHCC quasi-maximum likelihood estimates.

The table reports the parameter estimates and their asymptotic standard deviations in parenthesis. The log-likelihood value, the Akaike information criterion, and the Bayesian information criterion are reported at the bottom of the table.

Specification		all	1 yr	2 yrs	3 yrs	4 yrs	5 yrs	7 yrs	10 yrs
AC(2)	RMSE	4.37	4.59	3.23	3.99	5.08	4.79	2.71	5.55
	Vol	4.37	4.58	2.79	3.55	4.98	4.41	2.69	5.55
	Min	-17.80	-8.35	-10.49	-16.30	-17.80	-8.54	-9.44	-15.14
	Max	28.57	28.57	22.20	18.67	13.39	15.41	8.38	20.18
	Med	-0.30	-0.87	-0.92	-0.92	0.88	2.37	-0.06	-1.04
AC(3)	RMSE	2.04	1.64	2.05	1.98	1.92	2.60	2.13	1.79
	Vol	2.03	1.56	2.00	1.81	1.79	2.19	1.97	1.78
	Min	-15.14	-5.22	-10.22	-15.03	-15.14	-5.55	-7.92	-10.40
	Max	10.04	9.76	10.04	4.46	3.09	9.50	5.55	6.81
	Med	-0.02	0.20	-0.17	-0.49	-0.32	1.07	-0.87	0.33

Table 4: Fitted spread errors for the affine models.

The table reports the minimal, maximal, median, and root mean squared errors by maturity over the entire time period for the two different specifications.

	AC(2)	AC(3)
κ_1	0.0183 (0.0066)	0.1209 (0.0764)
κ_2	0.0402 (0.0070)	0.0166 (0.0115)
κ_3		0.0646 (0.0650)
θ_1	0.0000 (0.0406)	0.0001 (0.0365)
θ_2	0.5999 (0.0263)	0.2630 (0.1381)
θ_3		0.0506 (0.5301)
σ_1	0.0800 (0.0066)	0.1222 (0.8818)
σ_2	0.3497 (0.0784)	0.2962 (0.3224)
σ_3		0.5998 (0.6035)
X_{10}	0.0024 (0.0005)	0.0074 (0.1216)
X_{20}	0.0489 (0.0117)	0.0142 (0.1088)
X_{30}		0.0291 (0.1927)
σ_{cds}	4.7509 (0.0000)	2.3359 (0.0000)
r	0.0253	0.0253
$\log\mathcal{L}$	23,131	25,298
AIC	-46,243	-50,571
BIC	-46,187	-50,489

Table 5: Quasi-maximum likelihood estimates for the affine models. The table reports the parameter estimates and their asymptotic standard deviations in parenthesis. The log-likelihood value, the Akaike information criterion, and the Bayesian information criterion are reported at the bottom of the table.




Figure 1: State space of the LHC model with a single risk factor, $m = 1$, and illustrations of the inward pointing drift conditions at the boundaries.




Figure 2: Drift and diffusion of the default intensity for the one-factor linear model (black line) and affine model (grey line). The parameter values are $\ell_1 = 0.03$, $\ell_2 = 0.20$, $\gamma = 0.10$, and $\sigma = 0.20$.




Figure 3: CDS data.

The left figure displays the CDS spreads in basis points on *JPMorgan Chase & Co* with maturities 1 year (black), 5 years (grey), and 10 years (light-grey). The right figure displays the weekly changes in 1-year versus 5-year CDS spreads.




Figure 4: Fit and errors.

The first row displays the fitted CDS spreads in basis points with maturities 1 year (black), 5 years (grey), and 10 years (light-grey) for the three specifications. The second row displays the root-mean-square error (in basis points) computed every day and aggregated over all the maturities.




Figure 5: Fitted factors.

The filtered factors of the three estimated specifications are displayed over time. The first row displays the drift only survival process, the second row the implied default intensity, and the last row the process X_{mt} in black and the process X_{2t} in grey for the three-factor models.

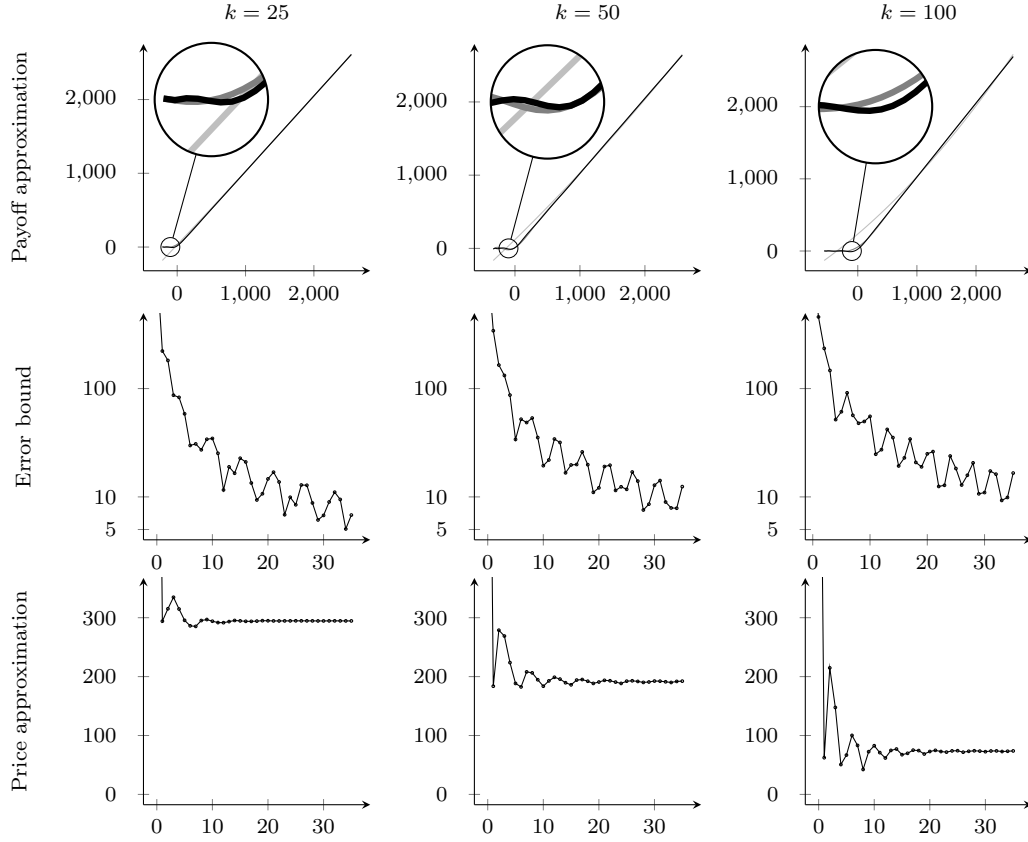


Figure 6: Payoff and price approximations of CDS options.

The first row displays the polynomial interpolation of the payoff function approximation by means of Chebyshev polynomials of order 2 (light-grey), 10 (grey), and 20 (black). The second and third rows display the price error bound and price approximation, respectively, as functions of the polynomial interpolation order. The first (second and third) column corresponds to a CDS option with a strike spread of 25 (50 and 100) basis points. All values are reported in basis points.




Figure 7: Fitted factors for the affine models.
The filtered factors of the two estimated affine specifications are displayed over time. The first row displays the implied default intensity X_{1t} , and the second row the process X_{mt} in black and the process X_{2t} in grey for the three-factor model.

References

- Ackerer, D., D. Filipović, and S. Pulido (2016). The Jacobi stochastic volatility model. Technical report, Swiss Finance Institute.
- Al-Mohy, A. H. and N. J. Higham (2011). Computing the action of the matrix exponential, with an application to exponential integrators. *SIAM Journal on Scientific Computing* 33(2), 488–511.
- Bielecki, T. R., M. Jeanblanc, and M. Rutkowski (2006). Hedging of credit derivatives in models with totally unexpected default. *Stochastic Processes and Applications to Mathematical Finance*, 35–100.
- Bielecki, T. R., M. Jeanblanc, and M. Rutkowski (2011). Hedging of a credit default swaption in the CIR default intensity model. *Finance and Stochastics* 15(3), 541–572.
- Bielecki, T. R., M. Jeanblanc, M. Rutkowski, et al. (2008). Pricing and trading credit default swaps in a hazard process model. *The Annals of Applied Probability* 18(6), 2495–2529.
- Bielecki, T. R. and M. Rutkowski (2002). *Credit risk: modeling, valuation and hedging*. Springer Science & Business Media.
- Brigo, D. and A. Alfonsi (2005). Credit default swap calibration and derivatives pricing with the SSRD stochastic intensity model. *Finance and Stochastics* 9(1), 29–42.
- Brigo, D. and N. El-Bachir (2010). An exact formula for default swaptions’ pricing in the SSRD stochastic intensity model. *Mathematical Finance* 20(3), 365–382.
- Brigo, D. and M. Morini (2005). CDS market formulas and models. In *Proceedings of the 18th annual Warwick options conference*.
- Cheridito, P., D. Filipović, and R. L. Kimmel (2007). Market price of risk specifications for affine models: Theory and evidence. *Journal of Financial Economics* 83(1), 123 – 170.
- Cheridito, P., D. Filipović, and M. Yor (2005). Equivalent and absolutely continuous measure changes for jump-diffusion processes. *The Annals of Applied Probability* 15(3), 1713–1732.
- Christoffersen, P., C. Dorion, K. Jacobs, and L. Karoui (2014). Nonlinear Kalman filtering in affine term structure models. *Management Science* 60(9), 2248–2268.

- Corrado, C. J. and T. Su (1996). Skewness and kurtosis in S&P500 index returns implied by option prices. *Journal of Financial Research* 19(2), 175–192.
- Delbaen, F. and H. Shirakawa (2002). An interest rate model with upper and lower bounds. *Asia-Pacific Financial Markets* 9(3-4), 191–209.
- Duarte, J. (2004, October). Evaluating an alternative risk preference in affine term structure models. *Review of Financial Studies* 17(2), 379–404.
- Duffee, G. (2002). Term premia and interest rate forecasts in affine models. *The Journal of Finance* LVII(1), 405–443.
- Duffie, D., J. Pan, and K. Singleton (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68(6), 1343–1376.
- Duffie, D. and K. Singleton (1999). Modeling term structures of defaultable bonds. *Review of Financial Studies* 12(4), 687–720.
- Elliott, R. J., M. Jeanblanc, and M. Yor (2000). On models of default risk. *Mathematical Finance* 10(2), 179–195.
- Filipović, D. and M. Larsson (2016). Polynomial diffusions and applications in finance. *Finance and Stochastics*. Forthcoming.
- Filipović, D., M. Larsson, and A. Trolle (2016). Linear-rational term structure models. *Journal of Finance*. Forthcoming.
- Filipović, D., E. Mayerhofer, and P. Schneider (2013). Density approximations for multivariate affine jump-diffusion processes. *Journal of Econometrics* 176(2), 93–111.
- Gabaix, X. (2009). Linearity-generating processes: A modelling tool yielding closed forms for asset prices. Working Paper 13430, NBER.
- Gaß, M., K. Glau, M. Mahlstedt, and M. Mair (2015). Chebyshev interpolation for parametric option pricing.
- Gourieroux, C. and J. Jasiak (2006). Multivariate Jacobi process with application to smooth transitions. *Journal of Econometrics* 131(1), 475–505.
- Higham, N. J. (2008). *Functions of matrices: theory and computation*. Siam.

- Hull, J. and A. White (1987). The pricing of options on assets with stochastic volatilities. *The Journal of Finance* 42(2), 281–300.
- Hull, J. C. and A. D. White (2003). The valuation of credit default swap options. *The Journal of Derivatives* 10(3), 40–50.
- Jamshidian, F. (2004). Valuation of credit default swaps and swaptions. *Finance and Stochastics* 8(3), 343–371.
- Jarrow, R. and A. Rudd (1982). Approximate option valuation for arbitrary stochastic processes. *Journal of Financial Economics* 10(3), 347–369.
- Karatzas, I. and S. E. Shreve (1991). *Brownian motion and stochastic calculus* (Second ed.), Volume 113 of *Graduate Texts in Mathematics*. New York: Springer-Verlag.
- Lando, D. (1998). On Cox processes and credit risky securities. *Review of Derivatives research* 120, 99–120.
- Rudin, W. (1974). *Functional Analysis*. McGraw-Hill.
- Schönbucher, P. (2000). A Libor market model with default risk.
- Schönbucher, P. (2004, August). A measure of survival. *Risk*, 79–85.
- Sidje, R. B. (1998). Expokit: A software package for computing matrix exponentials. *ACM Transactions on Mathematical Software (TOMS)* 24(1), 130–156.